



Skull segmentation from MR scans using a higher-order shape model based on convolutional restricted Boltzmann machines

Puonti, Oula; Van Leemput, Koen; Nielsen, Jesper Duemose; Bauer, Christian; Siebner, Hartwig Roman; Madsen, Kristoffer Hougaard; Thielscher, Axel

Published in:
Medical Imaging 2018: Image Processing

Link to article, DOI:
[10.1117/12.2293073](https://doi.org/10.1117/12.2293073)

Publication date:
2018

Document Version
Publisher's PDF, also known as Version of record

[Link back to DTU Orbit](#)

Citation (APA):
Puonti, O., Van Leemput, K., Nielsen, J. D., Bauer, C., Siebner, H. R., Madsen, K. H., & Thielscher, A. (2018). Skull segmentation from MR scans using a higher-order shape model based on convolutional restricted Boltzmann machines. In E. D. A., & B. A. L. (Eds.), *Medical Imaging 2018: Image Processing* (Vol. 10574). SPIE - International Society for Optical Engineering. <https://doi.org/10.1117/12.2293073>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

PROCEEDINGS OF SPIE

[SPIDigitalLibrary.org/conference-proceedings-of-spie](https://spiedigitallibrary.org/conference-proceedings-of-spie)

Skull segmentation from MR scans using a higher-order shape model based on convolutional restricted Boltzmann machines

Oula Puonti, Koen Van Leemput, Jesper D. Nielsen, Christian Bauer, Hartwig Roman Siebner, et al.

Oula Puonti, Koen Van Leemput, Jesper D. Nielsen, Christian Bauer, Hartwig Roman Siebner, Kristoffer H. Madsen, Axel Thielscher, "Skull segmentation from MR scans using a higher-order shape model based on convolutional restricted Boltzmann machines," Proc. SPIE 10574, Medical Imaging 2018: Image Processing, 1057404 (2 March 2018); doi: 10.1117/12.2293073

SPIE.

Event: SPIE Medical Imaging, 2018, Houston, Texas, United States

Skull segmentation from MR scans using a higher-order shape model based on convolutional restricted Boltzmann machines

Oula Puonti^{a,b}, Koen Van Leemput^{c,d}, Jesper D. Nielsen^{a,c}, Christian Bauer^a, Hartwig Roman Siebner^{a,e}, Kristoffer H. Madsen^{a,c}, and Axel Thielscher^{a,b}

^aDanish Research Center for Magnetic Resonance, Copenhagen University Hospital Hvidovre, Hvidovre, Denmark

^bDepartment of Electrical Engineering, Technical University of Denmark, Kgs. Lyngby, Denmark

^cDepartment of Applied Mathematics and Computer Science, Technical University of Denmark, Kgs. Lyngby, Denmark

^dDepartment of Radiology, Massachusetts General Hospital, Harvard Medical School, Charlestown, MA, USA

^eDepartment of Neurology, Copenhagen University Hospital Bispebjerg, Copenhagen, Denmark

ABSTRACT

Transcranial brain stimulation (TBS) techniques such as transcranial magnetic stimulation (TMS), transcranial direct current stimulation (tDCS) and others have seen a strong increase as tools in therapy and research within the last 20 years. In order to precisely target the stimulation, it is important to accurately model the individual head anatomy of a subject. Of particular importance is accurate reconstruction of the skull, as it has the strongest impact on the current pathways due to its low conductivity. Thus providing automated tools, which can reliably reconstruct the anatomy of the human head from magnetic resonance (MR) scans would be highly valuable for the application of transcranial stimulation methods. These head models can also be used to inform source localization methods such as EEG and MEG.

Automated segmentation of the skull from MR images is, however, challenging as the skull emits very little signal in MR. In order to avoid topological defects, such as holes in the segmentations, a strong model of the skull shape is needed. In this paper we propose a new shape model for skull segmentation based on the so-called convolutional restricted Boltzmann machines (cRBMs). Compared to traditionally used lower-order shape models, such as pair-wise Markov random fields (MRFs), the cRBMs model local shapes in larger spatial neighborhoods while still allowing for efficient inference. We compare the skull segmentation accuracy of our approach to two previously published methods and show significant improvement.

Keywords: skull segmentation, MRI, shape modeling, transcranial brain stimulation, head modeling

1. INTRODUCTION

Precise and reliable stimulation targeting in transcranial brain stimulation (TBS) is crucial both for studying the function of the brain in neuroscience research and for developing therapeutic intervention approaches for treatment of brain disorders such as stroke or depression. Accurate reconstruction of the individual head anatomy plays a crucial role in this, not only for locating the areas of interest for the stimulation, but also for modeling the effect of the individual anatomy on the field distribution inside the head which directly affects the spatial accuracy of the stimulation.¹ This is especially true for transcranial electric stimulation approaches, such as transcranial direct current stimulation (tDCS) and transcranial alternating current stimulation (tACS), where the skull has a large effect on the current flow due to its low conductivity,² and inaccuracies in the skull reconstruction can result in gross deviations in the location of the stimulus.³ Thus there is a clear need for automated tools that

Further author information: (Send correspondence to Oula Puonti)
Oula Puonti: E-mail: oulap@drcmr.dk

can accurately and robustly segment the different tissue classes, both brain and non-brain, directly from an MR scan of a given subject.

While there are many tools available for automated segmentation of different tissues and neuroanatomical structures in the brain, most of them do not provide segmentations of the extra-cranial tissues. The exceptions are the newest version of SPM (SPM12)⁴ and FSL⁵ which also provide the skull and skin segmentations but the field-of-view is limited to the upper part of the head with the neck excluded. Recently, Huang et al.⁶ proposed a fully automated segmentation framework which is specifically targeted for accurate reconstruction of the whole head anatomy with the neck included. Their framework relies on atlas-based segmentation (as implemented in SPM8), using an extended atlas of the whole head with six tissue classes:⁷ white matter, gray matter, cerebrospinal fluid, skull, skin and air. To avoid morphological errors and to ensure smooth segmentations, the authors apply a pairwise Markov random field (MRF) model, which encourages spatially clustered segmentations, as a post-processing step. The toolbox is called Morphologically and Anatomically Accurate Segmentation (MARS) and is provided as a part of the SPM8 software suite.

In this paper we propose a **higher-order** shape model for skull segmentation from MR images based on the convolutional restricted Boltzmann machine (cRBM) framework.^{8,9} Whereas the pairwise MRF model penalizes the number of differing label assignments typically between the six closest neighbors and the center voxel, the cRBM captures higher-order dependencies in the neighborhood of a given voxel, thus encoding shapes such as edges.¹⁰ Therefore, the cRBM model allows us to encode more realistic anatomical constraints, compared to the pairwise MRF, which in turn should provide better segmentation of the skull. In the experiments we compare the proposed shape model to the approach of Huang et al.⁶ and to the standard SPM post-processing using a simple Potts model.

2. METHODS

As the methods we compare against are implemented as post-processing steps to the segmentation function of SPM8, we apply the cRBM model in the same way to ensure fair comparisons. Therefore, in this section we first briefly describe the SPM8 segmentation framework, then the cRBM shape model and finally how the cRBM model can be combined with SPM8's segmentation model.

2.1 SPM8 – Unified Segmentation

Here we provide a short overview of the segmentation framework employed in SPM8, the detailed description can be found in Ashburner et al.⁴ Given an MR scan $\mathbf{d} = (d_1, \dots, d_I)$, where I denotes the number of voxels, and model parameters $\{\hat{\phi}, \hat{\nu}\}$, the task is to find a corresponding labeling $\mathbf{l} = (l_1, \dots, l_I)$ as the maximum-a-posteriori (MAP) estimate in a Bayesian framework:

$$\mathbf{l} = \arg \max_{\mathbf{l}} p(\mathbf{l} | \mathbf{d}, \hat{\phi}, \hat{\nu}) = \arg \max_{\mathbf{l}} p(\mathbf{d} | \mathbf{l}, \hat{\phi}) p(\mathbf{l} | \hat{\nu}), \quad (1)$$

where the first distribution on the right hand side is the likelihood, which models the MR data, and the second distribution is the prior, or atlas, which encodes the probabilities of different tissue classes in each voxel i . As mentioned in the previous section, the atlas we use has six classes and thus $l_i \in \{1, 2, \dots, k, \dots, 6\}$, i.e., the label in a voxel can take values from one to six.

The likelihood term models the MR intensities of each tissue k by a mixture of Gaussian distributions:⁴

$$p(\mathbf{d} | \mathbf{l}, \hat{\phi}) = \prod_i p(d_i | l_i, \hat{\phi}), \text{ where} \\ p(d_i | l_i = k, \hat{\phi}) = \sum_{g=1}^{\Lambda_k} \lambda_g \mathcal{N} \left(\rho_i(\hat{\beta}) d_i | \hat{\mu}_{k,g}, \hat{\sigma}_{k,g}^2 \right), \quad (2)$$

where Λ_k denotes the number of Gaussians for tissue k , $\mu_{k,g}$ is the mean, $\sigma_{k,g}^2$ is the variance and λ_g is the mixture weight. The scale term $\rho_i(\hat{\beta})$ refers to the so-called bias field, i.e., the low frequency noise corrupting MR images, which needs to be accounted for to obtain accurate segmentations.¹¹ The bias field is modeled as

a linear combination of the low-frequency components of the discrete cosine transform (DCT), where $\hat{\beta}$ are the coefficients. To shorthand the notation we let $\hat{\phi}$ denote all the parameters relating to the likelihood distribution, i.e., $\hat{\phi} = \{\hat{\mu}, \hat{\sigma}, \hat{\lambda}, \hat{\beta}\}$.

The prior distribution is defined as:⁴

$$p(\mathbf{l}|\hat{\nu}) = \prod_i p(l_i|\hat{\nu}), \text{ where}$$

$$p(l_i = k|\hat{\nu}) = \frac{\gamma_k b_{i,k}(\hat{\alpha})}{\sum_{k=1}^K \gamma_k b_{i,k}(\hat{\alpha})}. \quad (3)$$

Here γ_k denotes a global adjustment weight for tissue k , and $b_{i,k}(\hat{\alpha})$ denotes the probability of tissue k in voxel i where $\hat{\alpha}$ are the deformation parameters which map the atlas to the MR image. Similar to the bias field, the deformations are also modeled using DCT basis functions where $\hat{\alpha}$ are the coefficients. Again we let $\hat{\nu}$ denote all parameters relating to the prior distribution, i.e., $\hat{\nu} = \{\hat{\alpha}, \hat{\gamma}\}$.

Thus far, the parameters have been assumed to be known, which is naturally not the case when we are segmenting a target image. The parameter values are obtained as MAP estimates given the data:

$$\{\hat{\phi}, \hat{\nu}\} = \arg \max_{\phi, \nu} p(\phi, \nu | \mathbf{d}). \quad (4)$$

The optimization proceeds in two subsequent steps, where first the likelihood model parameters are updated keeping the prior parameters fixed and next the likelihood parameters are fixed while the prior parameters are updated. These two steps are repeated until the optimization has converged. Once the parameter estimates are found, a segmentation can be obtained from eq.(1) by assigning the most probable tissue label to each voxel separately as both the likelihood and prior distributions factorize over voxels.

2.2 cRBM – A Model of Shape

We model the shape of the skull using cRBMs^{8,9} which encode the shape of a binary (skull) image $\mathbf{z} = \{z_1 \dots z_I\}$ using a set of filters $\mathbf{w} = \{\mathbf{w}_1, \dots, \mathbf{w}_F\}$, where each filter is connected to a layer of binary hidden units $\mathbf{H} = \{\mathbf{h}_1, \dots, \mathbf{h}_F\}$, and F is the number of filters and corresponding groups of hidden units. The filters encode different features in a local neighborhood, whereas the hidden units act as feature detectors taking on a value of zero or one depending on whether a feature was detected in a given neighborhood. The cRBM defines a distribution over the image of the form:

$$p(\mathbf{z}) = \sum_{\mathbf{H}} p(\mathbf{z}, \mathbf{H}), \text{ where}$$

$$p(\mathbf{z}, \mathbf{H}) = \frac{1}{\mathcal{Z}} \exp(-E_{\text{cRBM}}(\mathbf{z}, \mathbf{H})). \quad (5)$$

Here the energy of the cRBM is defined as:⁸

$$E_{\text{cRBM}} = - \sum_f \mathbf{h}_f \bullet (\mathbf{w}_f * \mathbf{z}) - \sum_f c_f \sum_j h_{f,j} - \sum_i b_i z_i, \quad (6)$$

where \bullet denotes an element-wise product followed by summation and $*$ denotes a convolution. Furthermore, c_f is the so-called hidden bias parameter, which encourages the hidden units in group f to take on certain values, and b_i is the so-called visible bias parameter which does the same for each voxel in \mathbf{z} . Finally, \mathcal{Z} denotes the normalization term.

The visible bias is often assumed to be shared over the whole image, i.e., $b_i = b$. However, as we have access to the probabilistic atlas which encodes the probability for skull appearing in a given voxel we define $b_i = \log(p(l_i = \text{skull}|\hat{\nu}))$. The remaining parameters, i.e., $\theta = \{\mathbf{w}, \mathbf{c}\}$, are learned from a set of binary skull segmentations by performing stochastic gradient ascent on the log-likelihood using the so-called *contrastive divergence* approximation which has been shown to work well in practice.¹²

2.3 Incorporating the Shape Model into SPM8

To incorporate the shape model into SPM8's segmentation framework, we modify the segmentation prior to include the cRBM model:

$$p(\mathbf{l}|\hat{\boldsymbol{\nu}}, \hat{\boldsymbol{\theta}}) \propto \sum_{\mathbf{H}} p(\mathbf{l}, \mathbf{H}|\hat{\boldsymbol{\nu}}, \hat{\boldsymbol{\theta}}), \quad \text{where}$$

$$p(\mathbf{l}, \mathbf{H}|\hat{\boldsymbol{\nu}}, \hat{\boldsymbol{\theta}}) \propto \exp \left(\sum_i \log p(l_i|\hat{\boldsymbol{\nu}}) + \sum_f \hat{c}_f \sum_j h_{f,j} + \sum_f \mathbf{h}_f \bullet (\hat{\mathbf{w}}_f * \delta(\mathbf{l} = \text{skull})) \right). \quad (7)$$

In the lower formula we have used the fact that the visible bias in the cRBM is defined by the probabilistic prior, and in abuse of notation the first term refers back to eq.(3). The function $\delta(\cdot)$ returns one when the label in voxel i is skull and zero otherwise. The updated segmentation prior now consists of a unary term, which encodes the probability of a single voxel having one of the K labels, and a higher order term (second and third terms), which encode local skull shapes. Note that if the connections between the visible and hidden layers are disabled, i.e., the filters are set to zero, the model reverts back to the normal SPM8 model.

To obtain segmentations using this model, we first run the default SPM8 segmentation without the cRBM model to acquire initial estimates of the likelihood and atlas parameters. Due to the higher order terms in the segmentation prior, obtaining a MAP segmentation is analytically intractable as the prior distribution no longer factorizes over the voxels as in eq.(3). To obtain a segmentation we instead draw samples from $p(\mathbf{l}, \mathbf{H}, \phi|\mathbf{d}, \hat{\boldsymbol{\nu}}, \hat{\boldsymbol{\theta}})$ using blocked Gibbs sampling which is straight-forward, as each of the conditional distributions $p(\mathbf{l}|\mathbf{d}, \mathbf{H}, \phi, \hat{\boldsymbol{\nu}}, \hat{\boldsymbol{\theta}})$, $p(\mathbf{H}|\mathbf{l})$ and $p(\phi|\mathbf{d}, \mathbf{l})$ factorizes over its components. The final segmentation is obtained by voxel-wise majority voting over the collected samples.

3. EXPERIMENTS

3.1 Data

Our data set consists of 23 subjects who all underwent MRI and CT scanning. The MRI scans were performed on a 3T Philips Achieva MRI scanner using a 32-channel head coil. For each subject we acquired a high-resolution T1-weighted scan (repetition time = 6.0 ms; echo time = 2.7 ms; flip angle = 8 degrees; voxel size = 0.85mm^3) and a high-resolution T2-weighted scan (repetition time = 2500 ms; echo time = 272 ms; flip angle = 90 degrees; voxel size = 0.85mm^3). The low-dose CT scans were performed on a Siemens Biograph mCT (voxel size = $0.42 \times 0.42 \times 0.60\text{mm}$). The CT scans were coregistered to the MR scans and all scans were resampled to 1mm^3 resolution.

We used the CT scans to create ground truth skull segmentations by initially thresholding at 450 Hounsfield units, then applying a Chan-Vese level-set method¹³ to segment out the skull and finally using simple morphological operations to clean the segmentations. The segmentations were visually inspected to ensure good quality. The data set was then randomly divided into 13 training subjects and 10 test subjects. The training data set was used only for training the cRBM model and to cross-validate the free parameters for the methods we compare against. The test data set was used for comparing the performance of the different methods on unseen data.

3.2 Competing methods

We compare our approach to three different methods implemented in the SPM8 segmentation framework: the first method is the default SPM8 segmentation with no shape model post-processing (SPM8-default). The second method uses the standard SPM8 clean-up post-processing with a simple Potts model which encourages voxels to have the same label as their nearest neighbours (SPM-Potts). The third method uses the MARS post-processing which extends on the Potts model by modeling the co-occurrence of different label combinations in a neighborhood, as opposed to the simple Potts model which only encourages clustered segmentations with the same label.⁶ Thus MARS allows for encoding anatomical restrictions such as: *gray matter and skull labels can not occur together in a neighborhood*. Both the SPM-Potts and MARS have a single tunable parameter which relates to the strength of the spatial prior, i.e., how strongly the neighborhood configurations are enforced.

3.3 Implementation

Training the models We trained different cRBM models on the training data using 20 filters with sizes $(7 \times 7 \times 7)$, $(9 \times 9 \times 9)$ and $(11 \times 11 \times 11)$ to probe the effect of the neighborhood size on segmentation performance. Each of the cRBMs with different filter size was optimized using the contrastive divergence gradient approximation with 12000 gradient steps in total. The cRBM models were trained in SPM8 atlas space and the training data skull segmentations were non-linearly warped to the space of the atlas. Finally, after the models had been trained, we selected the optimal filter size by applying the different models on the 13 training subjects. Specifically, we computed a Dice score, which measures how well the automated skull segmentation corresponds to the ground truth, for all 13 subjects and picked the model that produced the best average Dice score over the training set. The best average was obtained using a filter size of $(11 \times 11 \times 11)$, although the difference to $(9 \times 9 \times 9)$ was quite small.

We followed a similar approach to tune the free parameters of the SPM-Potts and MARS post-processing approaches. Here we applied a grid search where the parameter value was varied as: $[1e-7, 1e-6, \dots, 0.1, 0.2, \dots, 1, 2, \dots, 10]$. For each parameter value we obtained skull segmentations of the 13 training subjects and computed the average Dice score against the CT-based ground truths. The best average score was obtained using parameter values of 5 and $1e-7$ for the Potts model and MARS respectively.

Applying the models As mentioned in section 2, analytical inference using the cRBM model is difficult and thus we resort to a sampling approach. This involves drawing samples in turn from the tissue segmentation \mathbf{I} , the hidden units of the cRBM model \mathbf{H} , and the Gaussian distribution parameters ϕ . After a burn-in period of 100 sampling steps, we collect 100 tissue segmentation samples and obtain the skull segmentation by voxel-wise majority voting over these samples. The initial segmentation and parameter estimates were obtained by running the default SPM8 segmentation method. Note that each of the different approaches uses the default SPM8 estimates as a starting point, and thus the only difference between the methods in the experiments is the post-processing step where different shape models are applied.

3.4 Results

The results are summarized in figure 1 and table 1. The proposed cRBM shape model outperforms both the Potts and MARS models when using only T1-weighted scans and when using a combination of T1-weighted and T2-weighted scans as input. The overall segmentation accuracy of all methods increases when a T2-weighted scan is included. This is probably because the border between skull and cerebro-spinal fluid (CSF) can be better delineated as CSF appears bright in T2-weighted scans. Note that for the MARS model, somewhat surprisingly, best results on the training data were achieved with a weight parameter of $1e-7$. This small weight value explains the similarity of the SPM8-default and MARS scores as the MARS shape prior is very weakly enforced. The default setting for the weight parameter in the MARS toolbox is 0.1, which was obtained by the authors by cross-validation on a single simulated subject from the BrainWeb repository.⁶ Thus it is possible that the default value is overfitted to that specific subject which could explain the rather large difference to the cross-validated value we found in this study.

Figure 2 shows example segmentations obtained using the different post-processing approaches along with the CT scan and CT-based ground truth segmentation. Looking at the differences between the shape models, pointed out by the green arrows in figure 2, it seems that both the Potts and cRBM models are good at removing spurious bone voxels and enforcing smoother segmentations, but the cRBM model tends to capture shape better as shown by the vertebrae segmentations. This is due to the fact that the Potts model only encourages a voxel to take the same label as its neighbors without weighing the neighbors differently, whereas the cRBM model allows for representing more complex shapes in the local neighborhood through the filters \mathbf{w}_f .

Although post-processing generally results in more accurate segmentations, there are still some inaccuracies especially in the neck. The vertebrae are particularly difficult to segment correctly as their contrast is quite similar to the surrounding tissues in the MR scans, and there is typically more non-linear deformations in the neck compared to the skull. Furthermore, the atlas we used in the experiments overestimates the bone in the frontal part of the head, especially around the paranasal sinuses (red arrow, figure 2). This global segmentation




Figure 1. Dice score box plots for different methods when using either only a T1-weighted scan (upper figure) or both T1-weighted and T2-weighted scans (lower figure). On each box, the central horizontal line is the median, the edges of the box extend to 25th and 75th percentiles, the whiskers extend out to 1.5 times the width of the box, and data points outside of that are considered outliers and denoted by gray points.

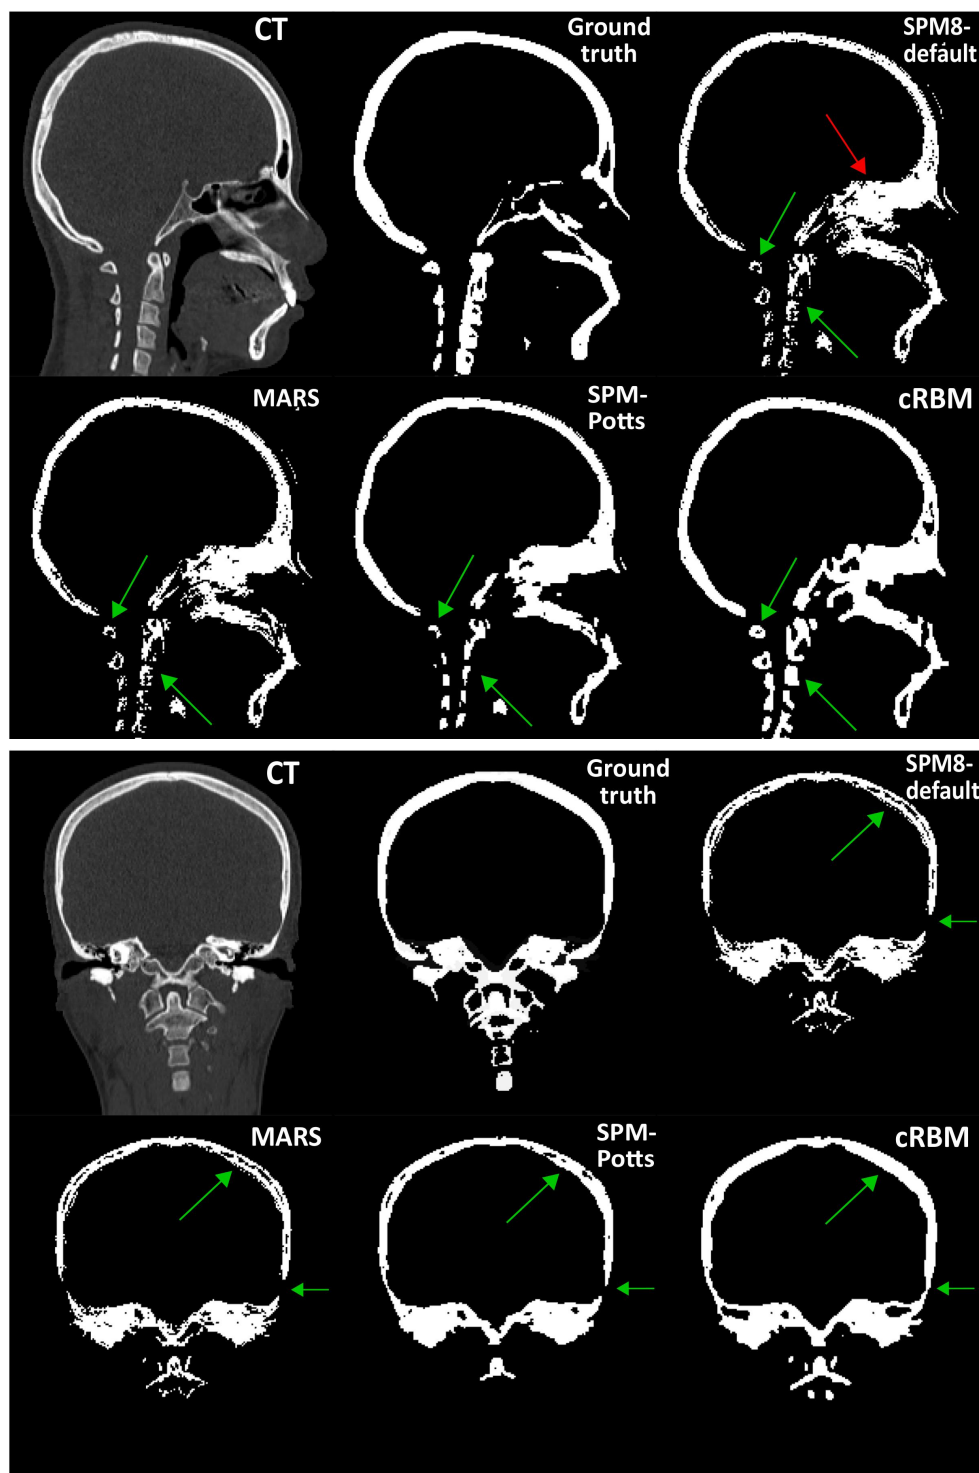


Figure 2. Example segmentations obtained using the different post-processing approaches. The upper figure shows a sagittal view and the lower one a coronal view. Both figures show (from left to right) the original CT scan, CT-based ground truth, default SPM8 segmentation with no post-processing, segmentation using MARS, segmentation using Potts model, and segmentation using cRBMs. The green arrows point to differences between the shape models, whereas the red arrow points to a global segmentation error (bone instead of air) resulting from the probabilistic atlas.

Methods/Contrasts	(1) SPM8-default	(2) SPM-Potts	(3) MARS	(4) cRBM
T1w	0.6780	0.6966 ^(1**,3**)	0.6780	0.7344 ^(1**,2**,3**)
T1w+T2w	0.7052	0.7194	0.7052	0.7446 ^(1**,2**,3**)

Table 1. Mean scores for the different methods and contrast combinations: the best mean score is denoted in bold. Significant difference, measured using the Wilcoxon signed-rank test, between the different methods is indicated in the superscript where * denotes $p < 0.05$ and ** denotes $p < 0.01$.

error cannot be corrected using local shape models and would instead require a new, more accurate probabilistic atlas of the bone to be constructed. Finally, we note that the segmentation obtained from the MARS model is very similar to the SPM8-default segmentation, which is consistent with the Dice score results shown in figure 1.

4. CONCLUSIONS

In this paper we have proposed a higher-order shape model based on cRBMs for skull segmentation from MR scans. Compared to previous lower-order shape models, which mainly encourage smooth segmentations, the cRBM model is able to model more complicated local features such as edges. In the experiments we compared the segmentation accuracy of the proposed method to publicly available benchmark methods on 10 test subjects using Dice scores. The results showed that the cRBM model improved the segmentation accuracy significantly compared to the other post-processing approaches, both when using only a T1-weighted scan and when using a combination of T1- and T2-weighted scans as input. In the future we aim to validate the proposed approach on larger data sets, acquired using different scanners and scan settings, to probe the robustness and feasibility of the method.

ACKNOWLEDGMENTS

This study was supported by the Lundbeck foundation (grant R118-A11308 to Axel Thielscher and grant R59-A5399 - PI Hartwig Roman Siebner), the Novonordisk foundation (grant no. NNF14OC0011413) and a PhD stipend of the Sino-Danish Center to Jesper D. Nielsen.

REFERENCES

- [1] Windhoff, M., Opitz, A., and Thielscher, A., "Electric field calculations in brain stimulation based on finite elements: An optimized processing pipeline for the generation and usage of accurate individual head models," *Human Brain Mapping* **34**(4), 923–935 (2013).
- [2] Indahlastari, A., Chauhan, M., Schwartz, B., and Sadleir, R. J., "Changing head model extent affects finite element predictions of transcranial direct current stimulation distributions," *Journal of Neural Engineering* **13**(6), 066006 (2016).
- [3] Datta, A., Bikson, M., and Fregni, F., "Transcranial direct current stimulation in patients with skull defects and skull plates: High-resolution computational FEM study of factors altering cortical current flow," *NeuroImage* **52**(4), 1268–1278 (2010).
- [4] Ashburner, J. and Friston, K. J., "Unified segmentation," *NeuroImage* **26**(3), 839–851 (2005).
- [5] Smith, S. M., Jenkinson, M., Woolrich, M. W., Beckmann, C. F., Behrens, T. E. J., Johansen-Berg, H., Bannister, P. R., De Luca, M., Drobnjak, I., Flitney, D. E., Niazy, R. K., Saunders, J., Vickers, J., Zhang, Y., De Stefano, N., Brady, J. M., and Matthews, P. M., "Advances in functional and structural MR image analysis and implementation as FSL," *NeuroImage* **23**(SUPPL. 1), S208–S219 (2004).
- [6] Huang, Y. and Parra, L. C., "Fully automated whole-head segmentation with improved smoothness and continuity, with theory reviewed," *PLoS ONE* **10**(5) (2015).
- [7] Huang, Y., Dmochowski, J. P., Su, Y., Datta, A., Rorden, C., and Parra, L. C., "Automated MRI segmentation for individualized modeling of current flow in the human head," *Journal of Neural Engineering* **10**(6), 066004 (2013).

- [8] Lee, H., Grosse, R., Ranganath, R., Ng, A. Y., and Honglak Lee, R. G., “Unsupervised learning of hierarchical representations with convolutional deep belief networks,” *Communications of the ACM* **54**(10), 95–103 (2011).
- [9] Norouzi, M., Ranjbar, M., and Mori, G., “Stacks of convolutional restricted Boltzmann machines for shift-invariant feature learning,” in *[2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2009]*, 2735–2742 (2009).
- [10] Eslami, S. M., Heess, N., Williams, C. K., and Winn, J., “The shape Boltzmann machine: A strong model of object shape,” *International Journal of Computer Vision* **107**(2), 155–176 (2014).
- [11] Van Leemput, K., Maes, F., Vandermeulen, D., and Suetens, P., “Automated model-based bias field correction of MR images of the brain,” *IEEE Transactions on Medical Imaging* **18**(10), 885–896 (1999).
- [12] Hinton, G. E., “Training Products of Experts by Minimizing Contrastive Divergence,” *Neural Computation* **14**(8), 1771–1800 (2002).
- [13] Zhang, Y., Matuszewski, B. J., Shark, L.-K., and Moore, C. J., “Medical Image Segmentation Using New Hybrid Level-Set Method,” in *[2008 Fifth International Conference BioMedical Visualization: Information Visualization in Medical and Biomedical Informatics (MEDIVIS)]*, 71–76 (2008).